Atmospheric Environment 44 (2010) 4426–4431

Contents lists available at ScienceDirect

Atmospheric Environment

journal homepage: www.elsevier.com/locate/atmosenv

An adaptive reduction algorithm for efficient chemical calculations in global atmospheric chemistry models

Mauricio Santillana [a,*], Philippe Le Sager [b,1], Daniel J. Jacob [b], Michael P. Brenner [b]

[a] Harvard University Center for the Environment, Cambridge, MA 02138, United States
[b] School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, United States

ARTICLE INFO

Article history:
Received 13 April 2010
Received in revised form
16 July 2010
Accepted 21 July 2010

Keywords:
Atmospheric chemistry
Multi-scale analysis
Time-scale separation
Reduction of chemical kinetics

ABSTRACT

We present a computationally efficient adaptive method for calculating the time evolution of the concentrations of chemical species in global 3-D models of atmospheric chemistry. Our strategy consists of partitioning the computational domain into fast and slow regions for each chemical species at every time step. In each grid box, we group the fast species and solve for their concentration in a coupled fashion. Concentrations of the slow species are calculated using a simple semi-implicit formula. Separation of species between fast and slow is done on the fly based on their local production and loss rates. This allows for example to exclude short-lived volatile organic compounds (VOCs) and their oxidation products from chemical calculations in the remote troposphere where their concentrations are negligible, letting the simulation determine the exclusion domain and allowing species to drop out individually from the coupled chemical calculation as their production/loss rates decline. We applied our method to a 1-year simulation of global tropospheric ozone-NO_x-VOC-aerosol chemistry using the GEOS-Chem model. Results show a 50% improvement in computational performance for the chemical solver, with no significant added error.

1. Introduction

Understanding the global-scale dynamics of the chemical composition of our atmosphere is essential for addressing a wide range of environmental issues from air quality to climate change. Global 3-D models of tropospheric chemistry must solve a system of coupled non-linear advection-reaction partial differential equations representing the temporal evolution of the different reactive species. Typical chemical mechanisms include over 100 species with lifetimes ranging from milliseconds to many years, giving rise to very large and stiff systems of differential equations. Solving these equations is difficult and the development of efficient and accurate techniques to achieve this has inspired research for the past 40 years. See for example (Jacobson, 1999; Sportisse, 2007) and the multiple references therein. A particular challenge for global models is to describe the full range of chemical environments including concentrated source regions that may have a large global influence. We show here that it is possible to achieve substantial computational savings by an adaptive method that dynamically adjusts the chemical mechanism to be solved to the local environment. We demonstrate the accuracy and performance of the method by application to the GEOS-Chem global chemical transport model (Bey et al., 2001).

Tropospheric chemistry models simulate the chemical composition of the atmosphere using a set of coupled non-linear partial differential equations of the type:

$$\frac{\partial C_i}{\partial t} + \boldsymbol{u}\cdot\nabla C_i = P_i - L_i, \quad i = 1,\ldots,N \tag{1}$$

where C_i (x,t) represents the spatio-temporal evolution of the concentration of species i, $\boldsymbol{u}(\boldsymbol{x},t)$ is the wind velocity, $P_i = P_i(\{C_j\},\boldsymbol{x},t)$ is the ensemble of atmospheric sources, and $L_i = L_i(\{C_j\},\boldsymbol{x},t)$ is the ensemble of atmospheric sinks. The species coupling shows up locally in P_i and L_i, through the group of chemical species $\{C_j\}$ that produce or react with species i. P_i and L_i are also functions of the local radiative and meteorological environment. The number of species N is typically over 100.

Equations (1) are solved in 3-D models using operator splitting methods that integrate the advection and chemistry operators

* Corresponding author. Tel.: +1 617 495 5941.
E-mail address: msantill@fas.harvard.edu (M. Santillana).
[1] Now at Royal Netherlands Meteorological Institute.

1352-2310/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.atmosenv.2010.07.044

separately. This enormously reduces the degrees of freedom of the non-linear system (Hundsdorfer and Verwer, 2003). In a simulation including detailed oxidant chemistry, the solution of the stiff system of ordinary differential equations (ODE):

$$\frac{dC_i}{dt} = P_i(\{C_j\}) - L_i(\{C_j\}), \quad i = 1,\ldots,N, \tag{2}$$

is the most time consuming process in the chemistry operator integration. Approaches to speed up this process include fast computational algorithms (Jacobson and Turco, 1994), and efficient numerical schemes such as implicit Rosenbrock solvers (Sandu et al., 1997; Damian et al., 2002). Other or complementary strategies use asymptotic analysis arguments, parameterization techniques, species lumping, or simplifications of chemical processes in particular locations of the domain. Examples include the separation of fast and slow species (Young and Boris, 1977; Gong and Cho, 1993), the quasi-steady-state assumption (QSSA) methods (Hesstvedt et al., 1978), functional parameterization of box model results (Jacob et al., 1989), species lumping (Sportisse and Djouad, 2000), and the use of different mechanisms for different regions either with specified boundaries (Jacobson, 1995) or with locally determined boundaries (Rastigeyev et al., 2007). Most approaches based on asymptotic analysis or parameterizations have been optimized for a given type of chemical regime and do not have the flexibility for implementation in a global model with a very wide range of possible regimes.

Separation of chemical mechanisms by geographical domains is particularly attractive for global models. Much of the complexity in these models is driven by short-lived non-methane volatile organic compounds (NMVOCs) emitted at the surface, which oxidize to a cascade of decomposition products of varying lifetimes. These NMVOCs and their short-lived decomposition products can be neglected in most of the troposphere, and their location-dependent exclusion from the chemical mechanism can greatly speed up the chemical computation. A challenge is to formulate this exclusion properly. Pre-setting geographical boundaries (as in Jacobson, 1995) is problematic because of the continuum of atmospheric lifetimes in NMVOCs and their oxidation products, and because of fast processes (such as deep convection) that can occasionally cause short-lived species to influence the remote troposphere far from their point of emission. Rastigeyev et al. (2007) presented an algorithm to locally diagnose a "chemical boundary layer" for individual species outside of which the species would be excluded from the mechanism and its concentration extrapolated by exponential decay. But this was only implemented in idealized scenarios of simple flow and chemistry.

In this work, we propose an algorithm that partitions, locally and on the fly, the global domain into fast and slow regions for each species, and adapts the solution strategy to the local environment dynamically in space and time. In this way our algorithm is adaptive. The proposed algorithm is implemented and tested using realistic meteorological conditions and full ozone-NO_x-VOC-aerosol chemistry.

2. The algorithm

Let us first re-write the chemistry operator as

$$\frac{dC_i}{dt} = P_i - k_i C_i, \quad i = 1,\ldots,N \tag{3}$$

since generally (though not always) the loss terms L_i have first-order dependence on the species concentration C_i. Here k_i is an effective loss rate constant. Asymptotic analysis arguments such as those utilized in low-dimensional manifold reduction methods (Lowe and Tomlin, 2000; Kaper and Kaper, 2002) and reduced chemical models (Djouad and Sportisse, 2003) show that one can separate species, based on the relative magnitude of the right-hand side of (3), as fast if $dC_i/dt > \delta$ or slow if $dC_i/dt \leq \delta$, for a small parameter δ, and compute an approximate solution to the ODE system (3) by only solving for the the group of fast species as a coupled system, and solving for the slow species either as an algebraic system ($dC_i/dt \approx 0$) or using a low-order numerical scheme ($dC_i/dt \approx \delta$). By construction, this family of approximate solutions for different values of δ converges to the solution of the system (3) provided $\delta \to 0$. Following these ideas, in our algorithm, at every grid box and time step t_n, we identify as fast species those satisfying either $P_i(\boldsymbol{x}, t_n) > \delta$ or $L_i(\boldsymbol{x}, t_n) > \delta$, and calculate their concentrations $\{C_i^{n+1} = C_i\ (t_{n+1})\}$ at the next time step, $t_{n+1} = t_n + \Delta t$, in a coupled fashion using an efficient implicit ODE solver for stiff systems, such as Gear-type or Rosenbrock. The concentrations of the rest of the (slow) species will not change much over a given time step since $dC_i/dt \approx 0$ or at most $dC_i/dt \approx \delta$, suggesting that we can treat $P_i = P_i^*$ and $k_i = k_i^*$ as constant for $t \in [t_n, t_n + \Delta t]$ (Young and Boris, 1977; Hesstvedt et al., 1978) and analytically solve Eq. (3) to obtain the formula

$$C_i^{n+1} = \frac{P_i^*}{k_i^*} + \left(C_i^n - \frac{P_i^*}{k_i^*}\right) e^{-k_i^* \Delta t} \tag{4}$$

to calculate their evolution separately.

The choice of an appropriate small parameter or threshold δ can be guided by simple reasoning. The general problem of tropospheric oxidant chemistry is in large part driven by reactions of the OH radical and ensuing radical chains. OH has a daytime concentration $\sim 10^6$ molecules cm^{-3} and a lifetime ~ 1 s. It follows that the chemical production and loss rates for important species in the fast mechanism (i.e., the P_i and $k_i C_i$ terms) may be expected to be within a few orders of magnitude of 10^6 molecules cm^{-3} s^{-1}. A species for which both production and loss rates are less than 10^2 molecules cm^{-3} s^{-1}, is unlikely to play a significant role in coupling to other species in the fast mechanism. This does not necessarily mean that the species is atmospherically unimportant, only that it is not significantly coupled to the other species.

Even though non-conservation of mass can arise from solving fast and slow species using different expressions, these mass imbalances can be controlled by choosing an adequate threshold δ based on the production and loss rates. Indeed, in Sect. 4.2, we find that a long-duration simulation (1 year) using our adaptive algorithm for $\delta < 10^2$ molecules cm^{-3} s^{-1} compares successfully with a benchmark simulation, making the issue of mass conservation negligible.

3. Implementation

We implemented our algorithm using the GEOS-Chem model (version v8-02-02). GEOS-Chem is a state-of-the-art 3-D global model of tropospheric chemistry driven by assimilated meteorological observations from the Goddard Earth Observing System (GEOS) of the NASA Global Modeling and Assimilation Office (GMAO). The model simulates global tropospheric ozone-NO_x-VOC-aerosol chemistry. The full chemical mechanism for the troposphere involves 111 species and over 300 reactions.

The chemical mass balance equations are integrated using a Gear-type solver (SMVGEAR II, Jacobson (1995)). Stratospheric chemistry is not explicitly simulated and it is instead parameterized to provide a proper representation of cross-tropopause fluxes. We used GEOS-Chem with a $4^\circ \times 5^\circ$ horizontal resolution and 20 sigma levels in the vertical for the following examples. For a detailed description of the original model see (Bey et al., 2001).

We performed three sets of simulations. The first two consisted of one-week long simulations aimed at selecting an appropriate threshold δ, and initialized on July 1st, 2004. The third one consisted of one-year simulations for 2005 aimed at testing the algorithm for the range of relevant time-scales of the system. The initial condition for the latter was obtained by running the model during 6 months prior to Jan 1st, 2005 (From Jul 1st, 2004 to Jan 1st, 2005) using the default settings (including the standard solver: SMVGEAR II).

We implemented our algorithm using LSODES (Livermore Solver for Ordinary Differential Equations; Radhakrishnan and Hindmarsh (1993)) instead of SMVGEAR II because it could be more easily configured to solve a different non-linear ODE system at each grid box at each time step. We will elaborate on this issue in Sect. 4.

Fig. 1 illustrates the partitioning between fast and slow species for peroxyacetylnitrate (PAN) and isoprene for a threshold of $\delta = 10^2$ molecules cm^{-3} s^{-1} and for July 8, 2004 at 0 GMT. Isoprene is emitted by vegetation and has a chemical lifetime of $\sim$1 h against oxidation by OH. In continental boundary layers where isoprene emissions are important, one needs to track isoprene and its successive oxidation products (about 40 species in the GEOS-Chem mechanism; Horowitz et al. (1998)) as part of the ensemble of fast species. Outside of the continental boundary layer, however, isoprene concentrations drop to negligibly low levels and can then be excluded from solution in the fast mechanism (note localized exceptions over Brazil and Southeast Asia where deep convection injects isoprene to the upper troposphere). As air parcels further age, the isoprene oxidation products gradually drop out from the fast mechanism. In this manner, the fast mechanism is considerably simplified in the free troposphere or over the oceans to exclude most isoprene chemistry.

The right panels of Fig. 1 show the fast-slow boundary for PAN. PAN is a particularly complicated molecule in a scale-separation scheme. It is produced by photochemical oxidation of NMVOCs in the presence of NO_x, most vigorously in the continental boundary layer but also at a slower rate in the global troposphere from the long-lived VOCs ethane and acetone. Its lifetime against thermal decomposition is less than 1 h at room temperature but increases to several months in the upper troposphere. Its long lifetime at cold temperatures allows it to be transported on a global scale to eventually release NO_x thermally or photochemically in the remote troposphere. We see from Fig. 1 that PAN needs to be included in the fast mechanism everywhere except in places that are very cold and dark (note the boundary between fast/slow regions at the terminator line in the upper troposphere) or warm and very remote (Indian Ocean near the surface).

Fig. 2 shows the percentage of the 111 species in the GEOS-Chem chemical mechanism that are treated as fast, for a threshold $\delta = 10^2$ molecules cm^{-3} s^{-1}. The percentage of fast species globally is 39%. There is no grid box in the domain where all the 111 species are fast. This is because some species are fast only in the daytime while others are fast only at night. Fig. 3 shows, for selected species, the percentage of tropospheric grid boxes where they are fast. Short-lived hydrocarbons are fast in only a small fraction of the grid boxes. OH, ozone, and NO_2 are fast almost everywhere.

4. Testing the algorithm

We evaluated our adaptive algorithm by comparison with a benchmark LSODES simulation including all species treated as fast. We chose the Relative Root Mean Square (RRMS) metric as given by in Sandu et al. (1997)

$$d_{AB}(C_i) = \sqrt{\frac{1}{M}\sum_{\Omega}\left|\frac{C_i^A - C_i^B}{C_i^A}\right|^2} \tag{5}$$

where C_i^A and C_i^B are the concentrations of species i calculated in simulations A and B, respectively, Ω is the set of grid boxes where C_i^A

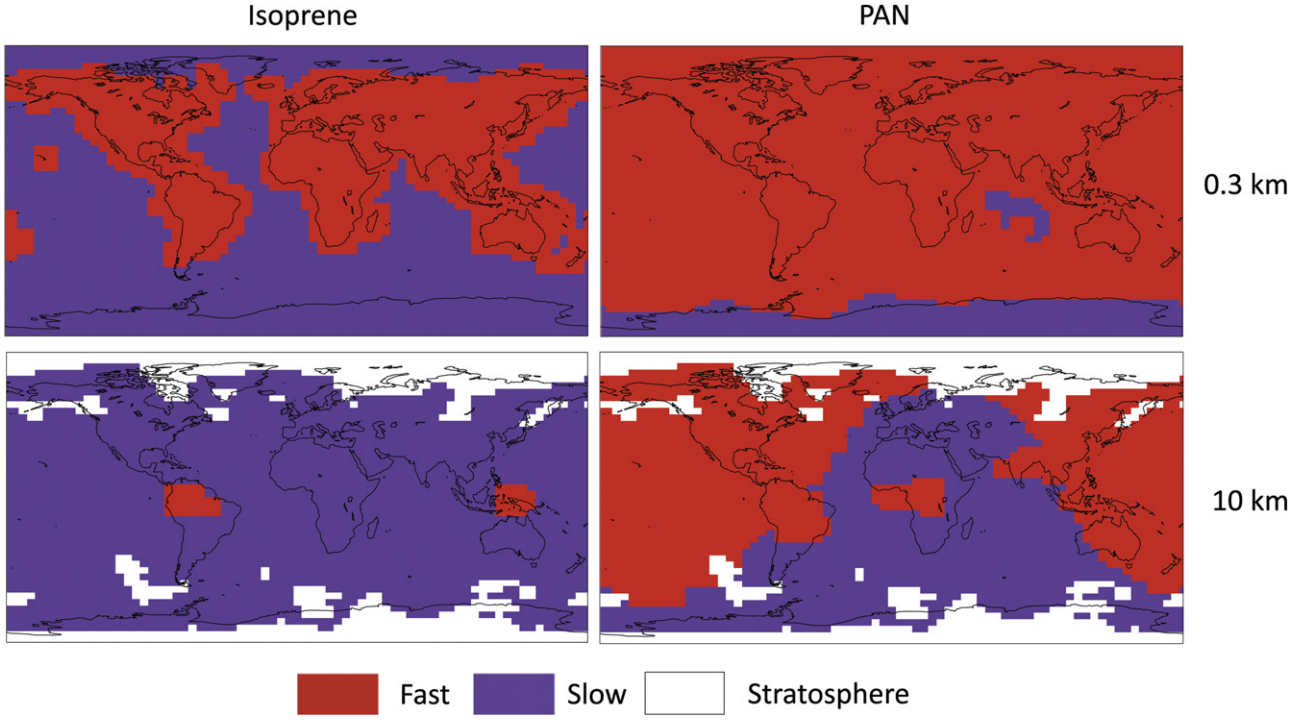

Fig. 1. Partitioning between fast and slow regions for isoprene and PAN (0.3 and 10 km above surface) using as threshold $\delta = 10^2$ molecules cm^{-3} s^{-1}. Results are from GEOS-Chem on July 8, 2004 at 00 GMT.

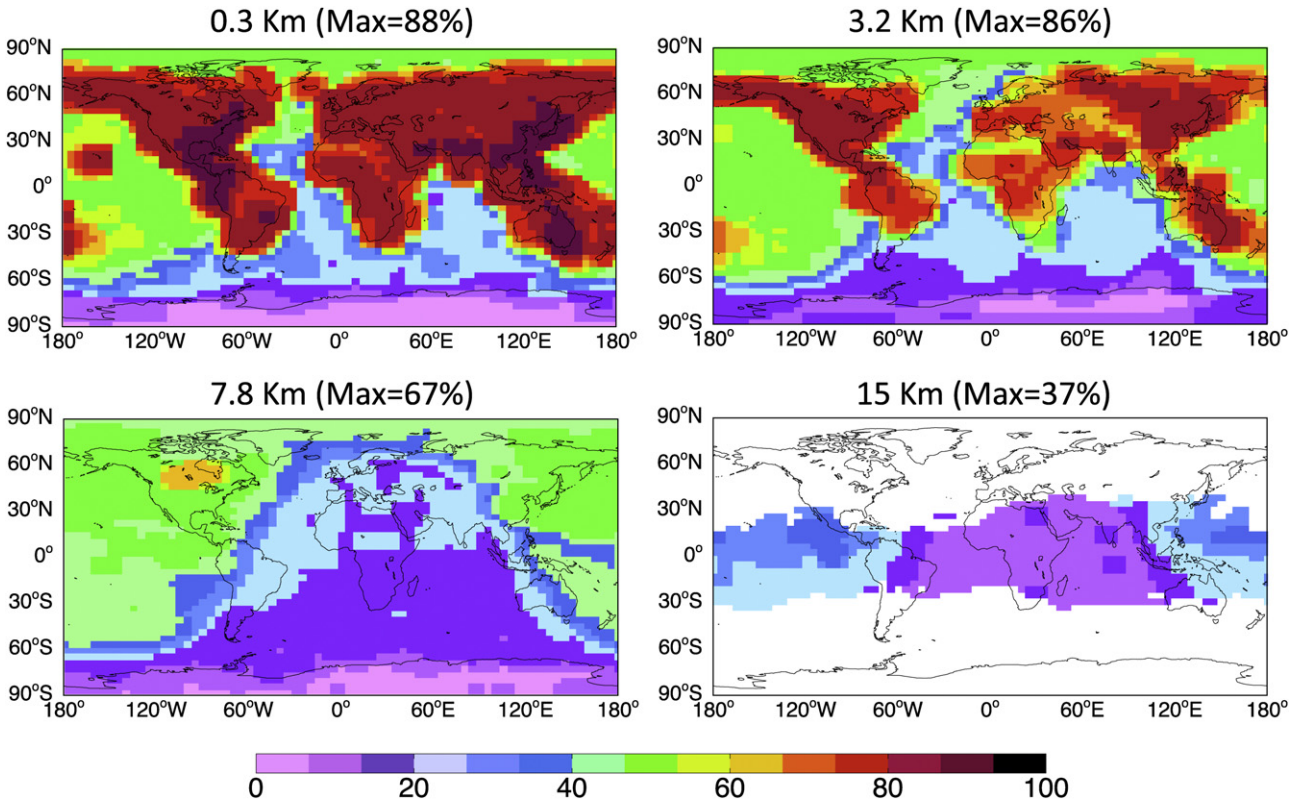

Fig. 2. Percentage of fast species in the GEOS-Chem chemical mechanism at different altitudes using a threshold of $\delta = 10^2$ molecules cm^{-3} s^{-1}. White boxes in the bottom right panel are in the stratosphere. Results are for July 8, 2004 at 00 GMT. The full GEOS-Chem chemical mechanism includes 111 species to describe tropospheric ozone-NOx-VOC-aerosol chemistry.

exceeds a threshold a, and M is the number of such grid boxes. We used $a = 10^6$ molecules cm^{-3} as in Eller et al. (2009) for our analysis.

4.1. Selection of threshold rate for fast mechanism

We performed two sets of 5 one-week GEOS-Chem ozone-NOx-VOC-aerosol chemistry simulations using the values $\delta = 0$ (all species are fast), 10, 10^2, 10^3, 10^4 molecules cm^{-3} s^{-1}. The first set of simulations were carried out in a chemistry-only mode (no transport) in order to assess the accuracy of our algorithm without the presence of any other mechanism. The second set was performed in a realistic configuration with all mechanisms of the code turned "on", we refer to this set as chemistry and transport simulations. In these numerical experiments, we compared the daily-averaged concentrations of species in the last day of the simulations using our adaptive algorithm for a given $\delta > 0$, with a benchmark simulation solving for all species as fast (standard method in GEOS-Chem).

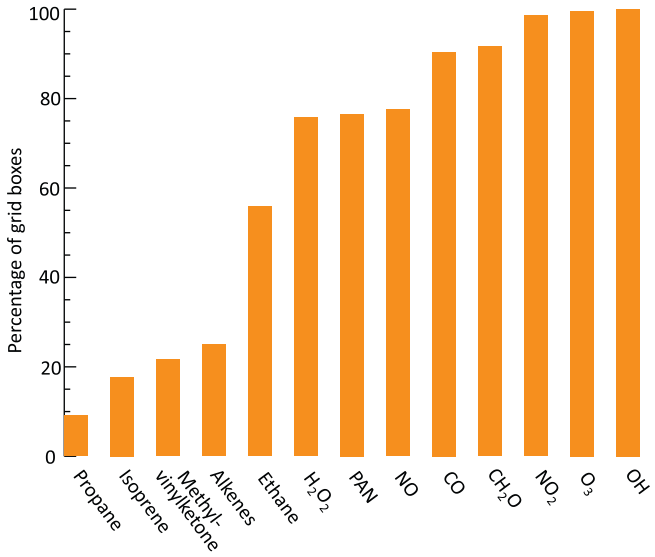

Fig. 3. Percentage of tropospheric grid boxes where particular species are fast using a threshold rate of $\delta = 10^2$ molecules cm^{-3} s^{-1}. Results are from a GEOS-Chem simulation on July 8, 2004 at 00 GMT.

The results of the experiments, presented in Fig. 4, show for each threshold δ an average value of the percentage of fast species during the simulation (top). The red data points show that the computational cost of the solver decreases linearly with the decrease of percentage of species, n, solved with the fast mechanism (100% corresponds to the CPU time used by the solver in a simulation with all species solved as fast, $\delta = 0$). This is a consequence of the $\mathcal{O}(n)$ efficiency of the sparse linear algebra routines used by LSODES. The blue and green data points show the median RRMS of the difference over all species between the benchmark simulation LSODES (all species fast) and the adaptive algorithm simulations for the different thresholds, in chemistry-only and chemistry and transport modes, respectively. Note that the RRMS of the differences of the chemistry and transport simulations are reduced when compared to the ones in chemistry-only mode. This may reflect the compensation of errors by transport, but also the maintenance of concentrations in a more realistic range over the 1-week simulation. In both modes, RRMS differences are negligible (<1%) when using a threshold $\delta = 10$ molecules cm^{-3} s^{-1} with significant savings of about 40%. For $\delta = 10^2$ molecules cm^{-3} s^{-1}, the median RRMS is less than 3% in the chemistry and transport

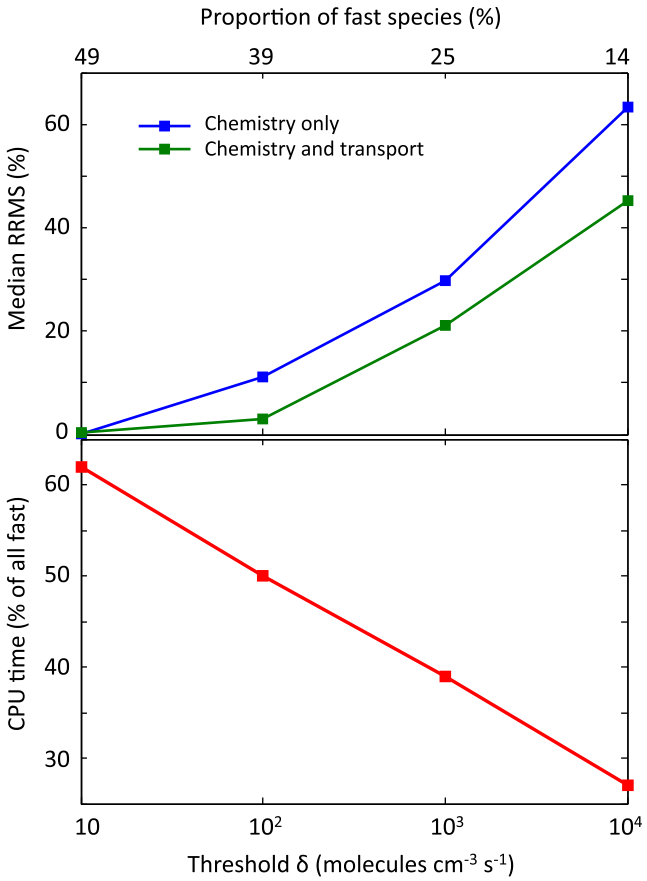

Fig. 4. Accuracy and performance of the adaptive reduction algorithm as a function of the production/loss rate threshold δ used to separate fast and slow species. The blue and green curves show the median RRMS over all species in chemistry-only and chemistry and transport modes respectively, the red curve shows the percentage of the computer time used for the chemical solver relative to a full-chemistry calculation ($\delta = 0$). The top scale shows the global percentage of species solved as fast for the different values of delta. Results are for 1-week simulations initialized on July 1, 2004. (For the interpretation of the reference to color in this figure legend the reader is referred to the web version of this article.)

simulations, while the computational time utilized by the solver is cut in half and the percentage of species being solved as fast is 39%. Using a higher threshold $\delta = 10^3$ molecules cm^{-3} s^{-1} incurs a significantly larger RRMS for only a 12% gain in the chemical solver computational requirement as compared to using $\delta = 10^2$ molecules cm^{-3} s^{-1}. Since the purpose of our algorithm is to provide an accurate solution in realistic simulations with coupling between chemistry and transport, we conclude that $\delta = 10^2$ molecules cm^{-3} s^{-1} is optimal.

4.2. One-year comparison

Based on the results of the previous section, a more comprehensive test of the accuracy of our algorithm was carried out for a one-year GEOS-Chem simulation of ozone-NOx-VOC-aerosol chemistry. This approach allows the sampling of the range of expected conditions and also tests the longer-lived species. We conducted three different one-year long simulations: the first one using our adaptive algorithm with a threshold value of $\delta = 10^2$ molecules cm^{-3} s^{-1}, the second one using LSODES with all species solved as fast, and the third one using GEOS-Chem's standard chemistry solver SMVGEAR II. All three simulations were identical except for the chemistry solver.

Fig. 5 shows the time evolution of the RRMS of differences between our algorithm and the standard (LSODES solver) simulations, for four selected species: NO_x, CO, OH, and O_3. Also shown is the comparison between the LSODES solver (all species fast) and SMVGEAR II. Note that the relative differences between our algorithm and the reference simulation do not exceed 5% for NO_x and are below 1% for CO, OH, and O_3. These differences are of the same magnitude as those observed between LSODES and SMVGEAR II with full-chemistry. We conclude that our adaptive algorithm does not induce significant errors relative to the full solution.

5. Conclusions

We presented a computationally efficient adaptive algorithm to calculate the time evolution of the concentrations of chemical species in global 3-D models of atmospheric chemistry.

Our algorithm identifies on the fly where a particular species is important (fast) or unimportant (slow) in the coupled chemical mechanism by comparing its production and loss rates with a prescribed threshold δ, and then adjusts the solution strategy accordingly. We solve for the concentration of fast species using the standard implicit solver of the model, and use an efficient semi-implicit formula for the slow species. The choice of the threshold δ is motivated by the observation that the characteristic magnitude of the production or loss rates driving the evolution of the coupled system are of the order of 10^6 molecules cm^{-3} s^{-1}. Thus, a species for which both production and loss rates are $< 10^2$ molecules cm^{-3} s^{-1} does not significantly drive the coupling of the system.

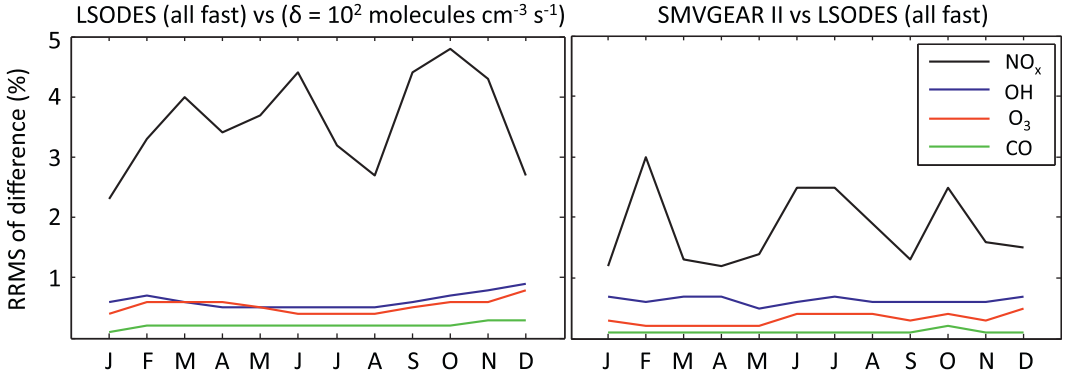

Fig. 5. Time evolution of the RRMS for four selected species (NO_x, CO, OH, and O_3) for one-year (2005) simulations. Daily average at the end of each month. The left panel compares a simulation using our algorithm with $\delta = 10^2$ molecules cm^{-3} s^{-1} to a reference simulation (LSODES $\delta = 0$, all species fast). The right panel compares the standard LSODES and SMVGEAR II solvers with all species fast.

A central advantage of our method for global 3-D model applications is to resolve the large differences in the complexity of the relevant mechanism between different regions of the world. For example, short-lived NMVOCs and their decomposition products may be important contributors to the chemical mechanism in continental boundary layers but not in the rest of the world. By diagnosing the relevant mechanism for a particular grid box on the fly, our method avoids the problematic prejudgment of where a species may be important.

We present a comprehensive evaluation of our algorithm in Sect. 4.1 with a 7-day ozone-NOx-VOC-aerosol chemistry simulation using different thresholds δ in the GEOS-Chem global 3-D model in a chemistry-only mode (transport turned off) and in the actual model with coupling between chemistry and transport. Coupling between chemistry and transport alleviates errors, which may reflect both error compensation during transport but also the maintenance of concentrations in more realistic ranges. In both cases, a threshold of $\delta = 10$ molecules cm^{-3} s^{-1} reduces the time integration of the chemistry solver by 40%, solving approximately 50% of species as fast, while showing negligible differences (<1%) when compared to a reference simulation. A larger threshold value of $\delta = 10^2$ molecules cm^{-3} s^{-1} still incurs errors of less than 3% in the coupled chemistry-transport simulation while cutting the time integration of the chemistry solver in half by solving approximately 40% of species as fast.

We show that the differences between a 1-year benchmark simulation, and a 1-year simulation using our algorithm do not grow in time, and in fact are of similar magnitude as differences between the two well-established solvers SMVGEAR and LSODES. These differences are less than 5% for NO_x and are below 1% for CO, OH, and O_3 during the year. This finding validates the efficiency of our approach for a broad range of time-scales and shows that issues of mass conservation are negligible for the choice $\delta = 10^2$ molecules cm^{-3} s^{-1}.

Acknowledgements

MS thanks the Harvard University Center for the Environment for the Henson Environmental Fellowship that funded his contribution for this investigation. PLS and DJJ were supported by the NASA Atmospheric Composition Modeling and Analysis Program. MPB was supported by the NSF Division of Mathematical Sciences. The authors thank Claire Carouge and Richard Ramaroson for valuable discussions. The authors thank the valuable comments and suggestions made by two anonymous reviewers.

References

Bey, I., Jacob, D.J., Yantosca, R.M., Logan, J.A., Field, B., Fiore, A.M., Li, Q., Liu, H., Mickley, L.J., Schultz, M., 2001. Global modeling of tropospheric chemistry with assimilated meteorology: model description and evaluation. Journal of Geophysical Research 106, 23,073–23,096.

Damian, V., Sandu, A., Damian, M., Potra, F., Carmichael, G.R., 2002. The kinetic preprocessor KPP-a software environment for solving chemical kinetics. Computers & Chemical Engineering 26 (11), 1567–1579.

Djouad, R., Sportisse, B., 2003. Solving reduced chemical models in air pollution modelling. Applied Numerical Mathematics 44 (1–2), 49–61.

Eller, P., Singh, K., Sandu, A., Bowman, K., Henze, D.K., Lee, M., 2009. Implementation and evaluation of an array of chemical solvers in the global chemical transport model GEOS-Chem. Geoscientific Model Development 2, 89–96.

Gong, W., Cho, H.R., 1993. A numerical scheme for the integration of the gas-phase chemical rate equations in three-dimensional atmospheric models. Atmospheric Environment 27A, 2147–2160.

Hesstvedt, E., Hov, O., Isaksen, I., 1978. Quasi-steady-state-approximation in air pollution modelling: comparison of two numerical schemes for oxidant prediction. International Journal of Chemical Kinetics 10, 971–994.

Horowitz, L.W., Liang, J., Gardner, G.M., Jacob, D.J., 1998. Export of reactive nitrogen from North America during summertime: sensitivity to hydrocarbon chemistry. Journal of Geophysical Research 103 (13), 451–476.

Hundsdorfer, W., Verwer, J.G., 2003. Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations. In: Springer Series in Computational Mathematics, 33. Springer.

Jacob, D.J., Sillman, S., Logan, J.A., Wofsy, S.C., 1989. Least-independent-variables method for simulations of tropospheric ozone. Journal of Geophysical Research 94, 8497–8509.

Jacobson, M.Z., 1995. Computation of global photochemistry with SMVGEAR-II. Atmospheric Environment 29 (18), 2541–2546.

Jacobson, M.Z., 1999. Fundamentals of Atmospheric Modeling. Cambridge University Press.

Jacobson, M.Z., Turco, R.P., 1994. SMVGEAR: a sparse-matrix, vectorized Gear code for atmospheric models. Atmospheric Environment 28A, 273–284.

Kaper, H.G., Kaper, T.J., 2002. Asymptotic analysis of two reduction methods for systems of chemical reactions. Physica D 165, 66–93.

Lowe, R., Tomlin, A., 2000. Low-dimensional manifolds and reduced chemical models for tropospheric chemistry simulations. Atmospheric Environment 34, 2425–2436.

Radhakrishnan, K., Hindmarsh, A.C. Description and Use of LSODE, the Livermore Solver for Ordinary Differential Equations, LLNL report UCRL-ID-113855, December 1993.

Rastigeyev, Y., Brenner, M.P., Jacob, D.J., 2007. Spatial reduction algorithm for atmospheric chemical transport models. Proceedings of the National Academy of Sciences 104, 13875–13880.

Sandu, A., Verwer, J.G., VanLoon, M., Carmichael, G.R., Potra, F.A., Dabdub, D., Seinfeld, J.H., 1997. Benchmarking stiff ODE solvers for atmospheric chemistry problems .1. Implicit vs explicit. Atmospheric Environment 31 (19), 3151–3166.

Sportisse, B., 2007. A review of current issues in air pollution modeling and simulation. Computational Geosciences 11, 159–181.

Sportisse, B., Djouad, R., 2000. Reduction of chemical kinetics in air pollution modeling. Journal of Computational Physics 164, 354–376.

Young, T.R., Boris, J.P., 1977. A numerical technique for solving stiff ordinary differential equations associated with the chemical kinetics of reactive flow problems. Journal of Physical Chemistry 81, 2424–2427.